

Oracle NoSQL Database Enterprise Edition, Version 18.1



Oracle NoSQL Database is a scalable, distributed NoSQL database, designed to provide highly reliable, flexible and available data management across a configurable set of storage nodes.

Data can be modeled as relational-database-style tables, JSON documents or key-value pairs. Oracle NoSQL Database is a sharded (shared-nothing) system which distributes the data uniformly across the multiple shards in the cluster, based on the hashed value of the primary key. Within each shard, storage nodes are replicated to ensure high availability, rapid failover in the event of a node failure and optimal load balancing of queries. NoSQL Database provides Java, C, C#, C++, Python and node.js drivers and a REST API to simplify application development. Oracle NoSQL Database is integrated with a wide variety of related Oracle and open source applications in order to simplify and streamline the development and deployment of modern big data applications. NoSQL Database is dual-licensed and available as an open-source Apache licensed Community Edition as well as a commercially licensed Enterprise Edition.

"Oracle NoSQL offers value to customers looking at ACID transactions; geodistributed data; application security with authentication and session-level SSL encryption, and integration with Oracle Database"

THE FORRESTER WAVE™:
BIG DATA NOSQL
Q3, 2016

KEY BUSINESS BENEFITS

- High Throughput
- Bounded latency
- Near-linear scalability
- High Availability
- Short time to deployment

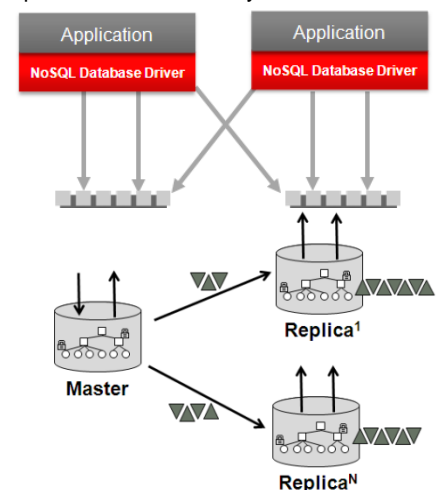
Architecture

The Oracle NoSQL Database is built upon the proven Oracle Berkeley DB Java Edition high-availability storage engine, which is in widespread use in enterprises across industries. In addition, it adds a layer of services for use in distributed environments.

The resulting solution provides distributed, highly available key/value storage that is well suited to large-volume, latency-sensitive applications.

High Availability and No Single Point of Failure

The Oracle NoSQL Database provides



- Smart Topology Management
- Online Elastic Configuration
- Enterprise grade software and support

KEY FEATURES

- Native JSON datatypes
- JSON datatype queries with SQL
- Java, C, C++ API's
- Fast, index Btree storage
- Dynamic Partitioning (sharding)
- Transparent load balancing
- Streaming Large Object Support
- Table based data modeling
- Secondary Index Support
- Streams support
- ACID compliant transactions
- Replication for HA, fault tolerance, fail-over, read scalability
- JMC and SNMP agents for system monitoring
- Online rolling upgrade
- Efficient Multi-Zone support
- Wire level data encryption using SSL
- Node level backup and restore
- Integrated with Apache Hadoop
- Secure Full Text Search
- Aggregation
- Parent Child Joins
- Zone Affinity

single-master, multi-replica database replication. Transactional data is delivered to all replica nodes with flexible durability policies per transaction. In the event the master replica node fails, a PAXOS-based automated fail-over election process minimizes downtime. This allows for scalability, fail-over, and hot-standby.

Transparent Load Balancing

The Oracle NoSQL Database Driver partitions the data in real time and evenly distributes it across the storage nodes. It is network topology and latency-aware, routing read and write operations to the most appropriate storage node in order to optimize load distribution and performance.

Oracle Database Integration via External Tables

Support for external table allows fetching from the Oracle NoSQL Database data from Oracle database using SQL statements such as Select, Select Count(*) etc. Once data from the Oracle NoSQL Database is exposed through external tables, one can access the data via standard JDBC drivers and/or visualize it through enterprise Business Intelligence tools.

Integration with Other Oracle Products

Oracle Event Processing (OEP) provides read access to the Oracle NoSQL Database via the Oracle NoSQL Database cartridge. Once the cartridge is configured, CQL queries can be used to query the data.

The Oracle Semantic Graph has developed a Jena Adapter for the Oracle NoSQL Database to store large volumes of RDF data (as triplets/quadruplets). This adapter enables fast access to graph data stored in the Oracle NoSQL Database via SPARQL queries. An integration with Oracle Coherence has been provided that allows the Oracle NoSQL Database to be used as a cache for Oracle Coherence applications, also allowing applications to directly access cached data from the Oracle NoSQL Database.

Streams Processing

Based on Reactive Streams, streams processing in Oracle NoSQL Database allow for a notification service that permits a user to subscribe to all logical changes (table row puts and deletes) made to an Oracle NoSQL Database store. Applications can be alerted to these changes which allows for asynchronous monitoring of database changes.

Oracle Big Data SQL and Hive Integration

Oracle Big Data SQL is a common SQL access layer to data stored in Hadoop, HDFS, Hive and the Oracle NoSQL Database. This allows customers to run queries on the Oracle NoSQL Database from a Hive or Oracle Database. Users can also run MapReduce jobs against data stored in the Oracle NoSQL Database that's configured for secure access. The latest release also supports both primitive and complex data types

JSON Query

Query your JSON datatypes with familiar SQL queries. This powerful feature gives developers the ability to use SQL to query schemaless JSON data. NoSQL now offers

USE CASES

- “Last mile” Big Data connectivity
- Click-through data capture
- High-Throughput event processing
- Social Network data capture
- Online retail customer view
- Mobile application back end infrastructure
- Real time sensor aggregation
- Network device monitoring and management
- Scalable authentication
- Content management
- Archiving

RELATED PRODUCTS

The following Oracle products are easily used in conjunction with Oracle NoSQL Database:

- Oracle Big Data Appliance
- Oracle Exadata
- Oracle Big Data SQL
- Oracle Berkeley DB
- Oracle SQL Developer
- Oracle Spatial and Graph

the flexibility of rich query over schemaless data along -side more structured queries. Oracle NoSQL Database is now a true multi-model database.

Full Text Search (FTS)

Gives users the ability to perform very secure fast text and indexed searches on data stored in Oracle NoSQL Database. FTS combines the TABLE interface with ElasticSearch (ES) for a powerful way to find documents that satisfy a query. This provides a high performant, secure full-text search of Tables stored in Oracle NoSQL Database. Connects to Elasticsearch using various security plugins (i.e. Shield, X-Pack, Kerberos/SPNEGO)

Time-To-Live

Allows for data to be stored for a specified period of time and then deleted automatically which is a critical requirement for sensor data capture in an Internet Of Things (IoT) service.

Arbiters

The ability to reduce hardware requirements using less replicas per shard instance.

Import/Export Capabilities

Data can be moved to and from Oracle NoSQL Database using a simple, data exchange format.

Server Side Filtering

The ability to access and process data in Oracle NoSQL Database from the Oracle Database. This feature dramatically improves performance.

Configurable Smart Topology

System administrators indicate how much capacity is available on a given storage node, allowing more capable storage nodes to host multiple replication nodes. Once the system knows about the capacity for the storage nodes in a configuration, it automatically allocates replication nodes intelligently. This results better load balancing for the system, better use of system resources and minimizing system impact in the event of storage node failure. Smart Topology also supports Data Centers, ensuring that a full set of replicas is initially allocated to each data center.

Elastic Configuration

The Oracle NoSQL Database includes a topology planning feature, with which an administrator can now modify the configuration of a NoSQL database, while the database is still online. This allows the administrator to:

- **Increase Data Distribution:** by increasing number of shards in the cluster, this increases the write throughput.
- **Increase Replication Factor:** by assigning additional replication nodes to each shard, which increases read throughput and system availability.
- **Rebalance Data Store:** by modifying the capacity of a storage node(s), the

system can be rebalanced, re-allocating replication nodes to the available storage nodes, as appropriate. The topology rebalance command allows the administrator to move replication nodes and/or partitions from over utilized nodes onto underutilized storage nodes or vice versa.

Easy Administration and Enhanced System Monitoring

The Oracle NoSQL Database provides an administration service, which can be accessed either from a web console or a command-line interface (CLI). This service supports core functionality such as the ability to configure, start, stop and monitor a storage node, without requiring manual effort with configuration files, shell scripts, or explicit database operations.

In addition it also allows Java Management Extensions (JMX) or Simple Network Management Protocol (SNMP) agents to be available for monitoring. This allows management clients to poll information about the status, performance metrics and operational parameters of the storage node and its managed services

Table Data Model

A tabular data structure is available, which simplifies application data modeling by leveraging existing schema design core concepts. Table model is layered on top of the distributed key-value structure, inheriting all its advantages and simplifying application design even further by enabling seamless integration with familiar SQL-based applications.

Secondary Index

Primary key only based indexing limits number of low latency access paths. Sometime application needs a few non-primary-key based paths to support the whole solution for the real-time system. Being able to define secondary index on any value field dramatically improves performance for queries.

APIs

The Oracle NoSQL Database includes Java and C APIs. These simple APIs allow the application developer to perform CRUD operations on the Oracle NoSQL Database. These libraries also include Avro support, so that developers can serialize key-value records and de-serialize key-value records interchangeably between C and Java applications.

Python and Node.JS drivers

The Oracle NoSQL Database has added support for NodeJS and Python drivers. Customers can now use their Python or NodeJS application to query, create, update or delete data that's stored in NoSQL using these drivers.

Oracle RESTful Services

The Oracle NoSQL Database includes support for exposing the Oracle NoSQL Database tables as REST API endpoints. The data stored in the tables can be created, queries, updated, and deleted using REST with JSON filters. This allows customers to build a REST-based application that can access data in either the Oracle Database or

the Oracle NoSQL Database.

Large Object Support

Stream based APIs are provided in the product to read and write Large Objects (LOBs) such as audio and video files, without having to materialize the value in its entirety in memory. This permits low latency operations across mixed workloads of objects of varying sizes.

Apache Hadoop Integration

KVAvroInputFormat and KVInputFormat classes are available to read data from the Oracle NoSQL Database natively into Hadoop Map/Reduce jobs. One use for this class is to read the Oracle NoSQL Database records into Oracle Loader for Hadoop.

High Performance

The Oracle NoSQL Database is network topology and latency aware. The Oracle NoSQL Database Driver working in conjunction with highly scalable, fault tolerant, high throughput storage engine enables a more granular distribution of resources and processing, which reduces the incidence of hot spots and provides greater performance on commodity based hardware.

Online Rolling Upgrade

Upgrade and patching is an important part of any software support cycle. The Oracle NoSQL Database provides facilities to perform a rolling upgrade, allowing a system administrator to upgrade all of the nodes in the Oracle NoSQL Database cluster while the database continues to remain online and available to clients.

Multi-Zone Deployment

The Oracle NoSQL Database supports the definition of multiple zones from within the topology deployment planner. It leverages the definition of these zones internally to intelligently allocate replication of processes and data, ensuring optimal reliability during hardware, network & power related failure scenarios.

There are two types of Zones: Primary zones contain nodes that can be served as masters or replicas and are typically connected by fast interconnects. Secondary zones contain nodes which can only be served as replicas. Secondary zones can be used to provide low latency read access to data at a distant location, or to offload read-only workloads, like analytics, report generation, and data exchange for improved workload management. The Oracle NoSQL Database allows users to continue business operations in the event of zone failures. This allows for any planned maintenance that results in the taking of one or more zones offline without impacting business operations.

Enterprise Security

OS- OS-independent, cluster-wide password-based user authentication and Oracle Wallet integration, enables greater protection from unauthorized access to sensitive data. Additionally, session-level Secure Sockets Layer (SSL) encryption and network port restrictions deliver greater protection from network intrusion. Oracle NoSQL Database now can use Kerberos integration for external authentication. This allows for

Oracle NoSQL Database to be easily integrated with customers' existing application which are already protected by Kerberos.

Commercial Grade Software and Support

The Oracle NoSQL Database overcomes a significant limitation faced by many enterprises considering the implementation of NoSQL databases—the need for full supportability. The Oracle NoSQL Database is a commercial product fully supported by Oracle. This gives organizations the confidence and reduces the risk they need to deploy the Oracle NoSQL Database in the production environments they depend on to manage their business-critical data.

Zone Affinity

Allows the administrator to place master nodes of a KVStore in primary zones that are in close network proximity to the user applications. This helps users to get predictable latency for write operations to the KVStore.

Parent-Child Joins

Oracle NoSQL Database includes support for a special kind of join among tables that belong to the same table hierarchy. This is implemented with a NESTED TABLES clause that is semantically equivalent to the LEFT-OUTER-JOINS defined by standard SQL and supported by all RDBMS implementations. The Left Outer Join creates a result set containing pairs of matching rows from the left and right tables and you would see a similar behavior in Oracle NoSQL.

Aggregation Functions


Aggregate functions in Oracle NoSQL Database iterate over the rows, evaluate an expression for each row, and aggregate the returned values into a single value. Syntactically Aggregate functions appear in the SELECT clause. Supported Aggregate functions supported are: sum, count, avg, min, max.

CONTACT US

For more information about Oracle NoSQL Database, visit oracle.com or call +1.800.ORACLE1 to speak to an Oracle representative.

ORACLE

CONNECT WITH US

-  blogs.oracle.com/nosql
-  facebook.com/oracle
-  twitter.com/oraclenosql
-  oracle.com

Integrated Cloud Applications & Platform Services

Copyright © 2018, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0116